

**SYSTEM AND METHOD FOR SECURING FIBER CHANNEL DRIVE ACCESS IN
A PARTITIONED DATA LIBRARY**

Peter Camble
56 Seymour Road
Bishopston, Bristol BS7 9HT
Citizenship: United Kingdom

Stephen Gold
Rock Cottage
Stone Lane
Winterbourne, Bristol BS36 1DJ
Citizenship: United Kingdom

Ian Crighton
1 Old Manor Cottage
Winterbourne Hill
Winterbourne, Bristol BS36 1JS
Citizenship: United Kingdom

RELATED APPLICATIONS

[0001] The present invention is related to the following copending and commonly assigned United States patent applications: serial number [30014510-1] entitled System and Method for Partitioning a Storage Area Network Associated Data Library, filed December 28, 2001; serial number [30014511-1] entitled System and Method for Partitioning a Storage Area Network Associated Data Library Employing Element Addresses, filed December 28, 2001; serial number [30014512-1] entitled System and Method for Managing Access To Multiple Devices in a Partitioned Data Library, filed December 28, 2001; serial number [30014513-1] entitled System and Method for Peripheral Device Virtual Functionality Overlay, filed December 28, 2001; serial number [30014514-1] entitled System and Method for Securing Drive Access to Media Based On Medium Identification Numbers, filed December 28, 2001; serial number [30014515-1] entitled System and Method for Securing Drive Access to Data Storage Media Based On Medium Identifiers, filed December 28, 2001; serial number [30014517-1] entitled Method for Using Partitioning to Provide Capacity on Demand in Data Libraries, filed December 28, 2001; serial number [30014518-1] entitled System and Method for Intermediating Communication with a Moveable Media Library Utilizing a Plurality of Partitions, filed December 28, 2001; and serial number [30008195-1], entitled System and Method for Managing a Moveable Media Library with Library Partitions,

filed December 28, 2001; the disclosures of which are hereby incorporated herein by reference.

TECHNICAL FIELD

[0002] The present invention generally relates to data storage and specifically to a system and method for securing fiber channel drive access in a partitioned data library.

BACKGROUND

[0003] In certain storage area networks (SAN) usage scenarios, such as may arise for storage service providers (SSPs), there are multiple customers attempting to share the same common SAN resources. In such cases, there is a need to ensure that customers can only see and access the storage resources they have been allocated and prevent them from accessing storage of other customers. For example, if a customer stores their critical business data with a SSP, then they generally do not want other customers of the SSP reading their data or even being aware that they have information stored with the SSP. The capability to partition a tape library is known. However, special hardware or special backup software as described below has been used to implement partitioning.

[0004] Existing software-based data library partitioning solutions typically employ a host system that restricts access to portions of a tape library. The host restrictions are implemented by a mediating software process on a host system to enforce partition restrictions. However, this approach is problematic. Specifically, the approach is undesirable if the data library is utilized in a SSP environment. In SSP environments, the data library and the host systems may belong to different entities (e.g., the SSP and the customers). Placement of software mediating processes on host systems is unattractive, because it increases the burden on the customers to make use of the storage service. Moreover, many customers are unwilling to allow other parties to place software on their host systems. Additionally, the software mediating process approach is typically incompatible with existing data back-up utilities, i.e., the software mediating process approach requires the use of specialized data back-up applications. Hence, users are effectively denied the ability to run desired backup software.

[0005] Existing fibre channel (FC) disk array firmware may be used to provide security in an FC redundant array of independent disks (RAID), since the disk array firmware has direct control over the array's ports connected to the SAN. Every host and device connection into the SAN generally has a unique FC-based world-wide-name (WWN), which can be used by an FC-based RAID to uniquely identify a device or host connection. Therefore, the FC-disk array firmware may be configured so that when a host attempts to send a small computer systems interface (SCSI) command to a FC-logical unit number (LUN) inside the RAID, the firmware will check the originating WWN from the server that sent the command against a list of authorized WWNs. If the WWN is on the list of authorized WWNs for that RAID FC-LUN, the SCSI command may be processed, if the WWN is not on the list of authorized WWNs for the RAID FC-LUN the command will be rejected. The list of authorized WWN's for each RAID FC-LUN may be configured via the existing management software for the RAID.

[0006] However, if a standard existing SCSI device, such as a data tape library is connected to a FC SAN via existing FC interfaces, such as existing FC tape drives in the library, it is not possible to secure these devices so that only certain hosts can access them, as individual existing FC tape drives do not support the FC WWN-based security discussed above. As a typical example, if a FC tape drive is connected to a SAN, it is visible to every server connected to that SAN. This circumstance is unacceptable for a SAN that offers secure storage resources to diverse customers. Existing solutions do not allow fibre channel tape drive devices to be secured in a SAN environment. The scheme to secure LUNs implemented in FC disk arrays, as discussed above, does not extend to securing physical tape drives that make up a logical partition within a SAN attached tape library.

[0007] FC switches have the capability of configuring security zones that define which WWNs or FC ports of a server can see which WWNs or FC ports of devices. However, this FC switch zoning does not extend to device LUNs, so it is only possible to provide security using such FC switch zoning at the FC port level. Even if tape libraries are directly attached on a FC SAN, it would be very difficult for a user to define security zones for the library tape drives. A data tape library can have multiple FC tape drives, and may be logically partitioned into partitions extending across multiple fibre channel tape drives. Therefore, it would be difficult for a user to correctly identify which FC ports and LUNs

should be associated together in the same security zone for an FC switch. Understandably, a user may easily make mistakes in such a manual configuration process.

[0008] Access to stand-alone native FC devices may be secured by using switch zoning, facilitated by a one-to-one relationship between a stand alone FC drive and an accessing user's WWN. In a data library, the library controller is typically placed behind a bridge. Configuring an FC switch for switch zoning to secure such a controller adds a process for a SAN administrator to implement and coordinate with users. FC switch configuration is not typically under control of a library's management card.

SUMMARY OF THE INVENTION

[0009] One embodiment of a storage area network associated data library partitioning system comprises a plurality of storage slot elements adapted to store data storage media, at least one set of at least one of the slots is assigned to one partition of a plurality of partitions, and a plurality of data transfer elements that are adapted to receive the media and transfer data to and from the media, each of at least one set of at least one of the data transfer elements is assigned to one of the partitions, at least one data transfer element of each of the partitions hosts a logical element designation of a virtual controller for each of the partitions, the virtual controllers restricting movement of the media to between the set of slots and the set of data transfer elements assigned to a same of the partitions.

[0010] A preferred embodiment of a method according to the present invention for partitioning a storage area network associated data library comprises establishing a plurality of partitions in the data library, each of the partitions comprising at least one storage slot element and at least one data transfer element, each of the slots adapted to store media, and each of the data transfer elements adapted to receive the media and transfer data to and from the media, assigning a different logical element designation to each of the library partitions and assigning a same logical element designation as a partition to a virtual controller hosted by at least one of the data transfer elements in last the partition, and restricting movement of the media to between the slots and the data transfer elements assigned to a same partition.

BRIEF DESCRIPTION OF THE DRAWING

[0011] FIGURE 1 is a diagrammatic illustration of a SAN operating in accordance with a preferred embodiment of the present invention; and

[0012] FIGURE 2 is a diagrammatic illustration of an example of a data library operating in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0013] The present invention is directed to system and method, which provide FC security for FC resources of a partitioned data library. A surrogate LUN for a library controller provided by one or more of the FC tape drives in an SCSI-based data library partitioning system and method may also be secured in accordance with the present invention. A physical data library implementing the present invention may be partitioned into multiple virtual library partitions, with each library partition having one or more physical drives, and a unique subset of library media slots, and a dedicated virtual library changer device LUN assigned to the partition as discussed below. Such a data library partitioning system and method is disclosed commonly-assigned in U.S. Patent application serial number [30014511-1] entitled "System and Method for Partitioning a Storage Area Network Associated Data Library Employing Element Addresses". Preferably the present invention does not require modification to existing library hardware for implementation. The present invention is preferably implemented employing firmware modifications to subject FC-based drives and library controller(s).

[0014] Turning to FIGURE 1, SAN 100 is shown. By way of example, first and second customer servers 101 and 102 are connected to SAN 100 via FC switch 103. RAID 104 may be partitioned using existing LUN-based RAID partitioning methods, for example, assigning first partition 105 to server 101 and second partition 106 to server 102. Zero downtime backups (ZDBs) may be performed of the data each server has on the RAID to tape library 108, via ZDB interconnectivity 107 between RAID 104 and tape library 108. Such ZDBs preferably employ data-mover firmware embodied in RAID 104 or other elements of SAN 100. Such ZDBs are preferably carried out without impinging on the processor operations or LAN capacity of servers 101 and 102. Tape library 108 is preferably partitioned employing the aforementioned system and method for library logical partitioning

to insure that data for server 101 is maintained in partition 109 separate from data for server 102, and that the data of server 102 is maintained in partition 110 separate from data for server 101. Such partitioning facilitates implementation of the security system and method of the present invention to ensure that the servers may not access each other's data even though their data is maintained in the same physical library.

[0015] Data tape library 200 employing a preferred embodiment of the present system and method is illustrated in FIGURE 2 as an example of a library that may be employed as library 108 of FIGURE 1. However, other library designs and/or capacities may embody the present system and method. Exemplar data tape library 200 has four FC tape drives 201-204 serving as data transfer elements; forty media storage slots 205 organized into four trays 206-209 of ten slots 205 each; two FC-to-SCSI bridges 210 and 211; a library management interface card or remote management card (RMC) 212; library controller 213 and robotic media transport 220. The bridges, drives, transport, RMC and controller are preferably interconnected by inter-integrated circuit bus (I²C) 214. Additionally, drives 201-204 and library controller 213 preferably communicate with each other using dedicated automated control interface (ACI) links 221- 224 or the like, independently extending between each drive 201-204 and controller 213. Preferably, each drive is a FC device and has a FC address on a SAN with which the library is associated.

[0016] For partitions employed by a preferred embodiment of the present system and method, a subset of media slots 205 and tape drives 201-204 should be assigned to each partition, and a virtual library controller or dedicated virtual library changer device should be addressable with respect to each partition for control of library robotic media transport 220. The example partitioning shown in FIGURE 2 is indicated by boxes 215, 216 and 217. As illustrated, SCSI LUN0 (230) corresponds to partition 215, SCSI LUN1 (231) corresponds to partition 216 and SCSI LUN2 (232) corresponds to partition 217. Mailslots or import/export elements may be assigned to each partition or configured for use by the entire library. Preferably, easily accessible media storage slots are configured as mailslots.

[0017] Preferably, a FC device in each partition, such as drives 201-204, may host one or more FC LUNs. SCSI commands to the drive itself are preferably directed to LUN 0. Each drive may present a virtual controller as surrogate LUN1. Preferably, only one drive in a partition presents a virtual controller for that partition. Controller 213 dictates which drive

in a partition presents the virtual controller. Controller 213 configures the drive to provide the virtual controller via ACI 221, 222, 223 or 224.

[0018] SCSI commands to a virtual controller LUN received by a drive are passed to controller 213 over the drive's ACI. Controller 213 sends SCSI responses back to the drive over the drive's ACI 214. The drive, in turn, sends these SCSI responses over the FC SAN from the virtual controller LUN. The SCSI commands and responses are preferably sent over the ACI in a suitable form, packaged as an ACI message packet. The drive's firmware preferably supports functionality to facilitate hosting the virtual controller or surrogate LUN and pass back and forth SCSI messages to and from controller 213 over the drive's ACI. It is irrelevant to a drive which partition it is in, nor is it pertinent to a drive which logical controller is being addressed by an SCSI command. Controller 213 determines and maintains which drive of a partition is hosting the logical controller LUN. So, since the ACI is a point-to-point connection, as opposed to a bus (i.e. there is an ACI port on the controller for each drive, each of which connects to only one drive), when controller 213 receives SCSI commands over an ACI link, the commands are addressed to one particular logical controller. Therefore, when controller 213 receives a SCSI command from a logical controller of a drive, controller 213 can identify the partition based on the originating drive.

[0019] For each partition configured there will be one drive that hosts the logical controller LUN for that partition. As indicated above, the drive hosting the logical controller for the partition is determined by controller 213. Advantageously, if a drive in a partition fails, or is inadvertently disconnected from the FC SAN, the controller may configure one of the other drives in the partition to take over the logical library LUN hosting for that partition.

[0020] Access to existing stand-alone native FC devices may be restricted by using switch zoning, as discussed above. This is facilitated by the one-to-one relationship between an existing stand-alone FC drive and an accessing user's WWN. However, in a partitioned SCSI data library, library controller 213 is preferably placed behind a bridge, such as FC-to-SCSI bridge(s) 210 and/or 211. In such a situation, configuring an FC switch for switch zoning to secure controller 213 adds a process for a SAN administrator to implement and coordinate with users. FC switch configuration is not traditionally under control of a library management card and manual configuration of switch zoning is prone to error.

[0021] In accordance with the present inventive system and method native FC tape drives 201-204 may support security based on WWN or other unique host device identifiers without the need for switch zoning and the related manual configuration. To provide a more usable one-step configuration process such security may be established and modified via management card 212.

[0022] If all the tape drives 201 through 204 deployed in library 200 are FC tape drives and library controller 213 is not on a common bus with an FC-to-SCSI bridge, such as bridges 210 or 211, the library can be configured so that an instance of the library controller, one per partition, is accessed as surrogate LUNs 230, 231 and 232, via one tape drive in each partition. In the example illustrated in FIGURE 2, surrogate LUN0 (230) for partition 215 is provided by drive 201 while surrogate LUN1 (231) and surrogate LUN2 (232) are provided by drives 203 and 204, respectively, for partitions 216 and 217, respectively. The FC security of tape drives 201-204 and library controller LUN(s) 230-232 is preferably configured by a user via RMC 212. Additionally, RMC 212 defines which tape drives are in which partition.

[0023] To provide security in this fibre channel environment, a user may also configure which SAN hosts have access to partition resources such as tape drives, library controller and media in each partition, via a control interface of RMC 212. This security configuration may be carried out via a web browser interface or via a network management protocol interface. For example, the user may select an active partition and configure the partition to either be unsecured, allowing all hosts access, or restrict access to a list of host WWNs or similar unique host device identifiers. To provide maximum flexibility, by default a partition's security level is preferably set to unsecured. To prevent all hosts from accessing a partition, the partition may be configured with an empty list of WWNs. Conversely, access by all hosts to disabled resources not in an active partition is preferably restricted.

[0024] Preferably, the security configuration of a tape drive applies to access to the tape drive itself, which will include any extended third-party copy command, such as ZDBs, that the tape drive supports. The security configuration of a tape drive will also preferably apply to any library controller surrogate LUN 230 through 232 the tape drive is hosting or supporting. Preferably, RMC 212 has no need to know which tape drive in a partition is hosting a surrogate LUN. Preferably, all tape drives in a partition have the same

security settings. Therefore, as long as one of the tape drives in a partition hosts a surrogate LUN, for example as shown for partition 215 of FIGURE 2, the surrogate LUN 230 and drives 201 and 202 under the surrogate LUN will have the required security settings applied. Preferably, as discussed above, the firmware of the library controller determines which tape drive holds the library controller surrogate LUN for that partition. Alternatively, the firmware of the controller and the firmware of the tape drives may negotiate as to which tape drive holds the library controller surrogate LUN for each partition.

[0025] Preferably, a FC drive blocks the ability for a host connected to the associated SAN to see the drive. In other words, the drive does not respond to any SCSI commands (e.g. SCSI inquiry, etc.) based on the host's WWN. However, because the WWN is not sent in each SCSI command, a drive preferably filters based on source ID for the host assigned by a name server, as detailed below.

[0026] When a partition is reconfigured, the FC security settings of a tape drive are preferably reconfigured. RMC 212 sends a security configuration request to library controller 213 over I²C bus 214. According to a preferred embodiment, library controller 213 passes the security configuration request, in the format of a special ACI command, to the tape drive(s) via the ACI port of the tape drive(s). Since the FC-LUN security in the tape drives is configured out-of-band via the ACI, the SCSI bus used to carry data to and from the drive need not be used to configure security.

[0027] FC commands generally do not contain the WWN of the originating host. However, FC commands use a source ID. Therefore, in accordance with the present invention a tape drive should also maintain an FC source ID-to-WWN mapping. The tape drive should gather information regarding source ID-to-WWN mappings from a SAN-associated name server at login, and issue a request state change notification to the name server to be informed of any changes in these mappings. If new WWNs are added to a security look-up table maintained by an FC tape drive, the drive should query the name server for the source ID of this new WWN. Preferably, the source ID of each incoming command, whether issued to a tape drive or a surrogate LUN hosted by an FC tape drive, will be compared against the FC drive or surrogate LUN's security configuration and used to determine security access. If the source ID matches the source ID mapped to a WWN in the

tabulated security settings then access is allowed. If the security setting for the drive or surrogate LUN is unsecured then access is allowed regardless of the source ID.

[0028] If security access to a partition is changed then the new security settings of that partition will preferably be sent to all tape drives in the partition. When a tape drive's firmware receives a security configuration request over the ACI it should erase its current security settings and then store in non-volatile random access memory (NVRAM) the new list of authorized WWNs, or an unsecured setting, contained in the security configuration request. A security configuration request to each affected FC tape drive may contain a list of authorized WWNs for that device. Where a library partition is unsecured and thus available to any initiator WWN, a security configuration request will leave a drive unsecured. The default configuration for a tape drive is preferably unsecured. Finally, if a security configuration request establishes an empty list of WWNs for a tape drive, the tape drive should not be part of an active partition and is thus disabled preferably disallowing only access at all to the drive by any user.

[0029] The library management firmware can use a security configuration request to clear any security information to an unsecured state. This may be required if the user wishes to set the library back to factory defaults or if the library management firmware detects a replacement FC tape drive that contains security information from another library which needs to be overwritten. If a tape drive is added or removed from a partition, the security settings of that tape drive are preferably altered to reflect the security settings of the new partition.

[0030] As noted above, preferably, only firmware modifications to an existing library are required to employ the present invention. The modifications may need to be made to tape drive firmware to implement surrogate LUN functionality and to implement WWN-based filtering. The firmware of the library controller may need to be modified to give the controller the ability to configure the FC drives to use multiple logical controller surrogate LUN functionality to configure the FC drives to use WWN based filtering on a per-partition basis. As pointed out above, preferably, no hardware modifications are required.

[0031] As one skilled in the art should recognize the present system and method is well-suited for use with other types of drive to SAN interfaces, for example internet small

computer systems interface (iSCSI). Preferably, the only change for iSCSI devices to use the present system and method is that the iSCSI equivalent of the FC source ID and/or WWN, such as iSCSI name, is used to authenticate initiators for access to secured devices.

10033010-123001